

Resource Logics with a Diminishing Resource

Natasha Alechina
University of Nottingham
Nottingham, UK
nza@cs.nott.ac.uk

Brian Logan
University of Nottingham
Nottingham, UK
bsl@cs.nott.ac.uk

ABSTRACT

Model-checking resource logics with production and consumption of resources is a computationally hard and often undecidable problem. We introduce a simple and realistic assumption that there is at least one *diminishing resource*, that is, a resource that cannot be produced and every action has a non-zero cost on this resource. An example of such resource is time. We show that, with this assumption, problems that are undecidable even for the underlying Alternating Time Temporal Logic, such as model-checking under imperfect information and perfect recall, become decidable for resource logics with a diminishing resource.

1 INTRODUCTION

There has been a considerable amount of work on multi-agent temporal logics interpreted over structures where agents' actions consume resources, or both produce and consume resources. Examples include an extension of Coalition Logic where actions consume resources and coalitional modalities are annotated with resource bounds ('agents in coalition A have a strategy of cost at most b to achieve ϕ ') (RBCL) [6, 8], a similar extension for Alternating Time Temporal Logic ATL (RB-ATL) [7], extensions of Computation Tree Logic and Alternating Time Temporal Logic with both consumption and production of resources (RTL, RAL) [11, 12], a variant of resource bounded ATL where all resources are convertible to money and the amount of money is bounded (PRB-ATL) [17, 18], an extension of PRB-ATL to μ -calculus [16], a version of ATL with more general numerical constraints (QATL*) [13], a version of RB-ATL where unbounded production of resources is allowed (RB \pm ATL) [2, 5]. The model-checking problem for such resource logics is decidable, though often not computationally tractable, when resources are only consumed or where the amount of resources is somehow bounded. [12]. For RAL with unbounded production of resources, the model-checking problem is undecidable, and this holds even for several of its fragments [12], although recently a fragment of RAL without the boundedness assumption has been found where the model-checking problem is decidable [3]. A slightly different semantics compared to RAL, but allowing unbounded production of resources, also results in a decidable model-checking problem for resource extensions of ATL such as RB \pm ATL [5]; the complexity of the model-checking problem for RB \pm ATL has been shown to be 2EXPTIME-complete in [2].

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

''
© 2018 Copyright held by the owner/author(s).

There exists also a large body of related work on reachability and non-termination problems in energy games and games on vector addition systems with state [10, 20]. In fact, complexity and decidability results for resource logics in [2] build on the results for single-sided vector addition systems with states [1, 14].

As far as we are aware, there is no work on model-checking resource logics under imperfect information. For ATL (without resources) under imperfect information and with perfect recall uniform strategies the problem is undecidable for three or more agents [19]. It is however decidable in the case of bounded strategies [23]. For two player energy games with imperfect information and a fixed initial credit the existence of a winning strategy is also decidable [15].

In this paper we consider a special kind of models for resource logics satisfying a restriction that one of the resources is always consumed by each action. It is a very natural setting which occurs in many verification problems for resource logics. The first obvious example of such a resource is time. Time is always 'consumed' by each action, and no agent in the system can turn back the clock and 'produce' time. When a verification problem has time as one of the explicit resource parameters, the restriction certainly applies. Other examples include systems where agents have a non-rechargeable battery and where all actions consume energy, e.g., nodes in a wireless sensor network; and systems where agents have a store of propellant that cannot be replenished during the course of a mission and all actions of interest involve manoeuvring, e.g., a constellation of satellites. We call this special resource that is consumed by all actions a *diminishing resource*.

From the technical point of view, the restriction to systems with a diminishing resource has the advantage that all strategies become bounded, even if for other resource types unbounded production is allowed. In the case of RB \pm ATL with a diminishing resource where the model-checking problem is already known to be decidable and 2EXPTIME-complete, we can produce simpler model-checking algorithms and a lower complexity bound (PSPACE if resource bounds are written in unary). In the case of RB \pm ATL with a diminishing resource under imperfect information, the result of [23] does not apply immediately because the bound is not fixed in advance, but the logic is indeed decidable and we get a new set of model-checking algorithms and a complexity bound. Finally, the decidability of RAL with a diminishing resource follows from the result on the decidability of RAL on bounded models [12], but the model-checking algorithms and the PSPACE upper bound (for resource endowments written in unary) are specific to RAL with diminishing resource and are new.

The rest of the paper is organised as follows. In Section 2, we introduce RB \pm ATL[#] with a diminishing resource, motivate changes to its syntax (we use the Release operator instead of 'Always' or

‘Globally’, and do not allow infinite resource bounds), give a model-checking algorithm and analyse its complexity. In Section 3 we introduce $\text{RB} \pm \text{ATL}_{iR}^\#$, which is $\text{RB} \pm \text{ATL}^\#$ under imperfect information and perfect recall, and give a model-checking algorithm for it and analyse its complexity. Finally in Section 4 we define RAL with diminishing resource, give a model-checking algorithm for it and show that the complexity is the same as for $\text{RB} \pm \text{ATL}^\#$.

2 $\text{RB} \pm \text{ATL}^\#$

The logic $\text{RB} \pm \text{ATL}$ was introduced in [4], and its model-checking complexity studied in more detail in [5] and [2]. Here we consider a variant of this logic without the idle action which is interpreted on finite paths. It contains a Release operator instead of Globally and does not allow infinite values in resource bounds. We use Release because it is not definable in ATL in terms of Next, Until and Globally [21] while Globally is definable in terms of Release, and it has a more intuitive meaning on finite computations.

As is the case with $\text{RB} \pm \text{ATL}$, the syntax of $\text{RB} \pm \text{ATL}^\#$ is defined relative to the following sets:

$\text{Agt} = \{a_1, \dots, a_n\}$ is a set of n agents, $\text{Res} = \{res_1, \dots, res_r\}$ is a set of r resource types, Π is a set of propositions, and $\mathcal{B} = \mathbb{N}^{\text{Res}^{\text{Agt}}}$ is a set of resource bounds (resource allocations to agents). Elements of \mathcal{B} are vectors of length n where each element is a vector of length r (the k th element of the i th vector is the allocation of the k th resource to the i th agent). We will denote by \mathcal{B}_A (for $A \subseteq \text{Agt}$) the set of possible resource allocations to agents in A .

Formulas of $\text{RB} \pm \text{ATL}^\#$ are defined by the following syntax

$$\phi, \psi ::= p \mid \neg\phi \mid \phi \vee \psi \mid \langle\langle A^b \rangle\rangle \phi \mid \langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi \mid \langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi$$

where $p \in \Pi$ is a proposition, $A \subseteq \text{Agt}$, and $b \in \mathcal{B}_A$ is a resource bound. Here, $\langle\langle A^b \rangle\rangle \phi$ means that a coalition A can ensure that the next state satisfies ϕ under resource bound b . $\langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$ means that A has a strategy to enforce ψ while maintaining the truth of ϕ , and the cost of this strategy is at most b . Finally, $\langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi$ means that A has a strategy to maintain ψ until and including the time when ϕ becomes true, or to maintain ψ forever if ϕ never becomes true, and the cost of this strategy is at most b .

The language is interpreted on resource-bounded concurrent game structures. Without loss of generality, we assume that the first resource type is diminishing, i.e., is consumed by every action.

Definition 2.1. A resource-bounded concurrent game structure with diminishing resource ($\text{RB-CGS}^\#$) is a tuple $M = (\text{Agt}, \text{Res}, S, \Pi, \pi, \text{Act}, d, c, \delta)$ where:

- Agt is a non-empty finite set of n agents,
- Res is a non-empty finite set of r resource types, where the first one is the distinguished diminishing resource
- S is a non-empty finite set of states;
- Π is a finite set of propositional variables and $\pi : \Pi \rightarrow \wp(S)$ is a truth assignment which associates each proposition in Π with a subset of states where it is true;
- Act is a non-empty set of actions
- $d : S \times \text{Agt} \rightarrow \wp(\text{Act}) \setminus \{\emptyset\}$ is a function which assigns to each $s \in S$ a non-empty set of actions available to each agent $a \in \text{Agt}$. We denote joint actions by all agents in Agt available at s by $D(s) = d(s, a_1) \times \dots \times d(s, a_n)$;

- $c : S \times \text{Act} \rightarrow \mathbb{Z}^r$ is a partial function which maps a state s and an action σ to a vector of integers, where the integer in position i indicates consumption or production of resource r_i by the action (negative value for consumption and positive value for production). We stipulate that the first position in the vector is always at most -1 (at least one unit of the diminishing resource is consumed by every action).
- $\delta : S \times \text{Act}^{|\text{Agt}|} \rightarrow S$ is a partial function that maps every $s \in S$ and joint action $\sigma \in D(s)$ to a state resulting from executing σ in s .

In what follows, we use the usual point-wise notation for vector comparison and addition. In particular, $(b_1, \dots, b_r) \leq (d_1, \dots, d_r)$ iff $b_i \leq d_i \forall i \in \{1, \dots, r\}$, $(b_1, \dots, b_r) = (d_1, \dots, d_r)$ iff $b_i = d_i \forall i \in \{1, \dots, r\}$, and $(b_1, \dots, b_r) + (d_1, \dots, d_r) = (b_1 + d_1, \dots, b_r + d_r)$. We define $(b_1, \dots, b_r) < (d_1, \dots, d_r)$ as $(b_1, \dots, b_r) \leq (d_1, \dots, d_r)$ and $(b_1, \dots, b_r) \neq (d_1, \dots, d_r)$. Given a function f returning a vector, we denote by f_i the function that returns the i -th component of the vector returned by f .

We denote by $\text{prod}(s, \sigma)$ the vector obtained by replacing negative values in $c(s, \sigma)$ by 0s: it is the vector of resources produced by action σ . We denote by $\text{cons}(s, \sigma)$ the vector obtained by first replacing positive values in $c(s, \sigma)$ by 0s and then replacing negative values by their absolute values: $\text{cons}(s, \sigma) = (|\min(0, c_1(s, \sigma))|, \dots, |\min(0, c_r(s, \sigma))|)$. It returns the positive costs on each resource of executing σ . In particular, $\text{cons}_i(s, \sigma) \geq 1$.

We denote the set of all finite non-empty sequences of states (finite computations) in a $\text{RB-CGS}^\# M$ by S^+ . We consider only finite computations because we are interested in computations possible under a finite resource bound, and in the presence of a diminishing resource which is required for any action, such computations are always finite. For a computation $\lambda = s_1 \dots s_k \in S^+$, we use the notation $\lambda[i] = s_i$ for $i \leq k$, $\lambda[i, j] = s_i \dots s_j \forall 1 \leq i \leq j \leq k$, and $|\lambda| = k$ for the length of λ .

Given a $\text{RB-CGS}^\# M$ and a state $s \in S$, a *joint action by a coalition* $A \subseteq \text{Agt}$ is a tuple $\sigma = (\sigma_a)_{a \in A}$ (where σ_a is the action that agent a executes as part of σ , the a th component of σ) such that $\sigma_a \in d(s, a)$. For a joint action σ by a coalition A , we denote by $\text{cons}(s, \sigma) = (\text{cons}(s, \sigma_a))_{a \in A}$ the vector of costs of the joint action, similarly for $\text{prod}(s, \sigma)$. The set of all joint actions for A at state s is denoted by $D_A(s)$.

Given a joint action by Agt $\sigma \in D(s)$, σ_A (a projection of σ on A) denotes the joint action executed by A as part of σ : $\sigma_A = (\sigma_a)_{a \in A}$. The set of all possible outcomes of a joint action $\sigma \in D_A(s)$ at state s is:

$$\text{out}(s, \sigma) = \{s' \in S \mid \exists \sigma' \in D(s) : \sigma = \sigma'_A \wedge s' = \delta(s, \sigma')\}$$

A *strategy for a coalition* $A \subseteq \text{Agt}$ in a $\text{RB-CGS}^\# M$ is a mapping $F_A : S^+ \rightarrow \text{Act}^{|\text{Agt}|}$ such that, for every $\lambda \in S^+$, $F_A(\lambda) \in D_A(\lambda[|\lambda|])$. A computation λ is consistent with a strategy F_A iff, for all i , $1 \leq i < |\lambda|$, $\lambda[i+1] \in \text{out}(\lambda[i], F_A(\lambda[1, i]))$. We denote by $\text{out}(s, F_A)$ the set of all computations λ starting from s that are consistent with F_A .

Given a bound $b \in \mathcal{B}_A$, a computation $\lambda \in \text{out}(s, F_A)$ is b -consistent with F_A iff, for every i , $1 \leq i < |\lambda|$,

$$\text{cons}(\lambda[i], F_A(\lambda[1, i])) \leq e_A(\lambda[i])$$

where $e_A(\lambda[i])$ is the amount of resources agents in A have in $\lambda[i]$: $e_A(\lambda[1]) = b$ and

$$e_A(\lambda[i+1]) = e_A(\lambda[i]) - \text{cons}(\lambda[i], F_A(\lambda[1, i])) + \text{prod}(\lambda[i], F_A(\lambda[1, i])).$$

In other words, the amount of resources any of the agents have is never negative for any resource type.

A computation λ is b -maximal for a strategy F_A if it cannot be extended further while remaining b -consistent (the next action prescribed by F_A would violate b -consistency).

The set of all maximal computations starting from state s that are b -consistent with F_A is denoted by $\text{out}(s, F_A, b)$. Note that this set is finite, the maximal length of each computation is bounded by b (or rather by the minimal value for any agent in A of b_{a1} : the bound on the first resource).

Given a RB-CGS[#] M and a state s of M , the truth of an RB \pm ATL[#] formula ϕ with respect to M and s is defined inductively on the structure of ϕ as follows:

- $M, s \models p$ iff $s \in \pi(p)$;
- $M, s \models \neg\phi$ iff $M, s \not\models \phi$;
- $M, s \models \phi \vee \psi$ iff $M, s \models \phi$ or $M, s \models \psi$;
- $M, s \models \langle\langle A^b \rangle\rangle \bigcirc \phi$ iff \exists strategy F_A such that for all b -maximal $\lambda \in \text{out}(s, F_A, b)$, $|\lambda| \geq 2$ and $M, \lambda[2] \models \phi$;
- $M, s \models \langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$ iff \exists strategy F_A such that for all b -maximal $\lambda \in \text{out}(s, F_A, b)$, $\exists i$ such that $1 \leq i \leq |\lambda|$, $M, \lambda[i] \models \psi$ and $M, \lambda[j] \models \phi$ for all $j \in \{1, \dots, i-1\}$.
- $M, s \models \langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi$ iff \exists strategy F_A such that for all b -maximal $\lambda \in \text{out}(s, F_A, b)$, either $\exists i$ such that $1 \leq i \leq |\lambda|$: $M, \lambda[i] \models \phi$ and $M, \lambda[j] \models \psi$ for all $j \in \{1, \dots, i\}$; or, $M, \lambda[j] \models \psi$ for all j such that $1 \leq j \leq |\lambda|$.

The most straightforward way of model-checking RB \pm ATL[#] is to adapt the model-checking algorithm for RB \pm ATL [5] and add a clause for $\langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi$. We present this simple algorithm here because we will use it in modified form in subsequent sections. It is however possible to do RB \pm ATL[#] model-checking more efficiently in the spirit of [17].

The algorithm is shown in Algorithm 1. Given a formula, ϕ_0 , we produce a set of subformulas $\text{Sub}(\phi_0)$ of ϕ_0 in the usual way. $\text{Sub}(\phi_0)$ is ordered in increasing order of complexity. We then proceed by cases. For all formulas in $\text{Sub}(\phi_0)$ apart from $\langle\langle A^b \rangle\rangle \bigcirc \phi$, $\langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$ and $\langle\langle A^b \rangle\rangle \square \phi$ we essentially run the standard ATL model-checking algorithm [9]. Labelling states with $\langle\langle A^b \rangle\rangle \bigcirc \phi$ makes use of a function $\text{Pre}(A, \rho, b)$ which, given a coalition A , a set $\rho \subseteq S$ and a bound b , returns a set of states s in which A has a joint action σ_A with $\text{cons}(s, \sigma_A) \leq b$ such that $\text{out}(s, \sigma_A) \subseteq \rho$. Labelling states with $\langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$ and $\langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi$ is more complex, and in the interests of readability we provide separate functions: UNTIL for $\langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$ formulas is shown in Algorithm 2, and RELEASE for $\langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi$ formulas is shown in Algorithm 3.

Both algorithms proceed by depth-first and-or search of M . We record information about the state of the search in a search tree of nodes. A *node* is a structure that consists of a state of M , the resources available to the agents A in that state (if any), and a finite path (sequence of nodes and edges) leading to this node from the root node. Edges in the tree correspond to joint actions by all agents. Note that the resources available to the agents in

a state s on a path constrain the edges from the corresponding node to be those actions σ_A where $\text{cons}(s, \sigma_A)$ is less than or equal to the available resources. For each node n in the tree, we have a function $s(n)$ that returns its state, $p(n)$ that returns the nodes on the path, $\text{act}(n)$ that returns the joint action taken to reach $s(n)$ from the preceding state on the path (i.e., the edge to n), and $e(n)$ that returns the vector of resource availabilities in $s(n)$ for A as a result of following $p(n)$. The functions $\text{act}_a(n)$ and $e_a(n)$ return the action performed by agent $a \in A$ in $\text{act}(n)$ and the resources available to agent a in $e(n)$ respectively. We use $p(n)[i]$ to denote the i -th node in the path $p(n)$, and $p(n)[1, j]$ to denote the prefix of $p(n)$ up to the j -th node. The function $\text{node}_0(s, b)$ returns the root node, i.e., a node n_0 such that $s(n_0) = s$, $p(n_0) = []$, $\text{act}(n_0) = \text{nil}$, and $e(n_0) = b$. The function $\text{node}(n, \sigma, s')$ returns a node n' where $s(n') = s'$, $p(n') = [p(n) \cdot n]$, $\text{act}(n') = \sigma$, and for all agents $a \in A$ $e_a(n') = e_a(n) + \text{prod}(s(n), \sigma_a) - \text{cons}(s(n), \sigma_a)$.

Algorithm 1 Labelling ϕ_0

```

1: function RB $\pm$ ATL#-LABEL( $M, \phi_0$ )
2:   for  $\phi' \in \text{Sub}(\phi_0)$  do
3:     case  $\phi' = p, \neg\phi, \phi \vee \psi$  standard, see [9]
4:     case  $\phi' = \langle\langle A^b \rangle\rangle \bigcirc \phi$ 
5:        $[\phi']_M \leftarrow \text{Pre}(A, [\phi]_M, b)$ 
6:     case  $\phi' = \langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$ 
7:        $[\phi']_M \leftarrow \{s \mid s \in S \wedge$ 
8:         UNTIL-STRATEGY( $\text{node}_0(s, b), \langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi)\}$ 
9:     case  $\phi' = \langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi$ 
10:       $[\phi']_M \leftarrow \{s \mid s \in S \wedge$ 
11:        RELEASE-STRATEGY( $\text{node}_0(s, b), \langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi)\}$ 
12:   return  $[\phi_0]_M$ 

```

Algorithm 2 Labelling $\langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$

```

1: function UNTIL-STRATEGY( $n, \langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$ )
2:   if  $s(n) \in [\psi]_M$  then
3:     return true
4:   if  $s(n) \notin [\phi]_M$  then
5:     return false
6:    $\text{Act}A \leftarrow \{\sigma \in D_A(s(n)) \mid \text{cons}(s(n), \sigma) \leq e(n)\}$ 
7:   for  $\sigma \in \text{Act}A$  do
8:      $O \leftarrow \text{out}(s(n), \sigma)$ 
9:      $\text{strat} \leftarrow \text{true}$ 
10:    for  $s' \in O$  do
11:       $\text{strat} \leftarrow \text{strat} \wedge$ 
12:        UNTIL-STRATEGY( $\text{node}(n, \sigma, s'), \langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$ )
13:    if  $\text{strat}$  then
14:      return true
15:   return false

```

When checking whether $\langle\langle A^b \rangle\rangle \psi_1 \mathcal{U} \psi_2$ or $\langle\langle A^b \rangle\rangle \psi_1 \mathcal{R} \psi_2$ is true in a state s , we examine paths whose length is bounded by the smallest resource bound b_{a1} on the first resource in b (since every action costs at least 1 unit of the first resource, any computation

Algorithm 3 Labelling $\langle\langle A^b \rangle\rangle\phi \mathcal{R} \psi$

```

1: function RELEASE-STRATEGY( $n, \langle\langle A^b \rangle\rangle\phi \mathcal{R} \psi$ )
2:   if  $s(n) \in [\psi]_M \cap [\phi]_M$  then
3:     return true
4:   if  $s(n) \in [\psi]_M \wedge$ 
5:      $\exists \sigma \in D_A(s(n)) : \text{cons}(s(n), \sigma) \not\leq e(n)$  then
6:     return true
7:   if  $s(n) \notin [\psi]_M$  then
8:     return false
9:    $ActA \leftarrow \{\sigma \in D_A(s(n)) \mid \text{cons}(s(n), \sigma) \leq e(n)\}$ 
10:  for  $\sigma \in ActA$  do
11:     $O \leftarrow \text{out}(s(n), \sigma)$ 
12:     $strat \leftarrow true$ 
13:    for  $s' \in O$  do
14:       $strat \leftarrow strat \wedge$ 
15:        RELEASE-STRATEGY( $node(n, \sigma, s'), \langle\langle A^b \rangle\rangle\phi \mathcal{R} \psi$ )
16:    if  $strat$  then
17:      return true
18:  return false

```

can contain at most b_{a1} steps). An over-approximation of the size of this search tree is $S^{\min_{a \in A}(b_{a1})}$.

LEMMA 2.2. *Algorithm 1 on input M, ϕ terminates after at most $O(|\phi| \times |M|^k)$ steps where k is the maximal value of the first resource bound in ϕ .*

LEMMA 2.3. *Algorithm 1 is correct.*

PROOF. The Boolean cases of the algorithm are standard.

The algorithm for $\langle\langle A^b \rangle\rangle\phi$ returns all states from where there is an action by A that costs less than b and all outcomes of this action satisfy ϕ . Essentially in each such state there is a one-step strategy satisfying $\langle\langle A^b \rangle\rangle\phi$. This is all we need because the rest of actions on this strategy can be arbitrary; the computations that are produced by the strategy do not need to satisfy any additional constraints apart from being maximal, i.e., eventually running out of resources (which they are guaranteed to do because of the first resource).

The algorithm for $\langle\langle A^b \rangle\rangle\phi \mathcal{U} \psi$ performs forward and-or search while making sure ϕ remains true, until ψ is reached. It returns true if and only if it finds a strategy where each computation reaches a ψ state before A run out of resources to carry on with the strategy, and ϕ holds along the computation up to the point ψ becomes true. Again actions after the ψ state can be arbitrary.

The algorithm for $\langle\langle A^b \rangle\rangle\phi \mathcal{R} \psi$ is similar to $\langle\langle A^b \rangle\rangle\phi \mathcal{U} \psi$ apart from two points. One is that the ψ state should also satisfy ϕ (the invariant holds not just on the path to a ϕ state but in the ϕ state itself). This is ensured by the test at line 2. The second difference is that there is another way to make $\langle\langle A^b \rangle\rangle\phi \mathcal{R} \psi$ true, which is to run out of resources while maintaining ψ . This is the reason for the test at line 4: if the invariant ψ is true in s and there is an action σ in $D_A(s)$ that would cause A to run out of resources, we return true because for this computation λ , the strategy F_A such that $F_A(\lambda) = \sigma$ ensures that λ is a b -maximal computation (and it satisfies ψ everywhere). \square

THEOREM 2.4. *The model-checking problem for $RB \pm ATL^\#$ is decidable in PSPACE (if resource bounds are written in unary).*

PROOF. From Lemmas 2.2 and 2.3 we have a model checking algorithm that solves the model checking problem for $RB \pm ATL^\#$. The complexity which results from the time bound in Lemma 2.2 can be improved by observing that the depth first search can be arranged using a stack and we only need to keep one branch at a time on the stack. The size of the stack is bounded by $\min_{a \in A}(b_{a1})$ and hence is polynomial if b is represented in unary. \square

3 $RB \pm ATL^\#$ WITH IMPERFECT INFORMATION AND PERFECT RECALL

Agents often have to act under imperfect information, for example, if states are only partially observable, an agent may be uncertain whether it is in state s or s' . This is represented in imperfect information models as a binary indistinguishability relation on the set of states for each agent a , \sim_a : if a cannot distinguish s from s' , we have $s \sim_a s'$. This relation can easily be lifted to finite sequences of states: if $s_1 \sim_a s'_1, s_2 \sim_a s'_2$, then $s_1 s_2 \sim_a s'_1 s'_2$. An essential requirement for strategies under imperfect information is that they are *uniform*: if agent a is uncertain whether the history so far is λ or λ' ($\lambda \sim_a \lambda'$), then the strategy for a should return the same action for both: $F_a(\lambda) = F_a(\lambda')$. Intuitively, the agent has no way of choosing different actions in indistinguishable situations. A strategy F_A for a group of agents A is uniform if it is uniform for every agent in A . In what follows, we consider *strongly uniform* strategies [22], which require that a strategy work from all initial states that are indistinguishable by some $a \in A$.

Unfortunately, model-checking for ATL under imperfect information with perfect recall uniform strategies, ATL_{iR} , is undecidable for more than three agents [19]. It is known that the model checking problem for ATL_{iR} with *bounded* strategies is decidable, while for *finite* strategies it is undecidable [23]. Bounded strategies are those that are defined for sequences of states of at most some fixed length k . In $RB \pm ATL_{iR}^\#$, there is no fixed bound on the size of strategies, since the size of strategy depends on the formula and the model. However, we can show that indeed the model checking problem for $RB \pm ATL_{iR}^\#$ with imperfect information and perfect recall strongly uniform strategies is decidable.

The model checking algorithms are similar to those given for $RB \pm ATL^\#$ in Section 2 in that they proceed by and-or depth first search, storing information about the state of the search in a search tree of nodes. However, in this case, the algorithms for Next, Until and Release also take a stack (list) of ‘open’ nodes B , a set of ‘closed’ nodes C in addition to an $RB \pm ATL^\#$ formula. B records the current state of the search while C records ‘successful’ branches (rather than all visited nodes). Uniformity is ensured if action choices are consistent with those taken after \sim_a sequences of states on all successful paths explored to date: $(n_1, \dots, n_k \sim_a n'_1, \dots, n'_k$ iff $s(n_1), \dots, s(n_k) \sim_a s(n'_1), \dots, s(n'_k)$). In addition, we assume functions $hd(u), tl(u)$ which return the head and tail of a list u , and $u \circ v$ which concatenates the lists u and v . (We abuse notation slightly, and treat sets as lists, e.g., use $hd(u)$ where u is a set, to return an arbitrary element of u , and use \circ between a set and a list.)

$M, s \models \langle\langle A^b \rangle\rangle \circ \phi$ under strong uniformity requires that there exists a uniform strategy F_A such that for all $a \in A$, if $s' \sim_a s$, then for all b -maximal $\lambda \in \text{out}(s', F_A, b)$: $|\lambda| > 1$ and $M, \lambda[2] \models \phi$. Similarly, in truth conditions for $\langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$ and $\langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi$ we require the existence of a uniform strategy where all b -maximal computations starting from states s' indistinguishable from s by any $a \in A$ satisfy the Until (respectively, Release) formula.

Weak uniformity only requires the existence of a uniform strategy from s . It is easy to modify the algorithms below to correspond to weak uniformity semantics. In fact, the algorithm for $\langle\langle A^b \rangle\rangle \circ \phi$ would become much simpler (identical to that for $\text{RB} \pm \text{ATL}^\#$ in the previous section).

Algorithm 4 Labelling ϕ_0

```

1: function  $\text{RB} \pm \text{ATL}^\#_{IR}\text{-LABEL}(M, \phi_0)$ 
2:   for  $\phi' \in \text{Sub}(\phi_0)$  do
3:     case  $\phi' = p, \neg\phi, \phi \vee \psi$  standard, see [9]
4:     case  $\phi' = \langle\langle A^b \rangle\rangle \circ \phi$ 
5:        $[\phi']_M \leftarrow \{s \mid s \in S \wedge$ 
           $\text{NEXT}([\text{node}_0(s', b) : s' \sim_{a \in A} s],$ 
           $\{ \}, \langle\langle A^b \rangle\rangle \circ \phi)\}$ 
6:     case  $\phi' = \langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$ 
7:        $[\phi']_M \leftarrow \{s \mid s \in S \wedge$ 
           $\text{UNTIL}([\text{node}_0(s', b) : s' \sim_{a \in A} s],$ 
           $\{ \}, \langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi)\}$ 
8:     case  $\phi' = \langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi$ 
9:        $[\phi']_M \leftarrow \{s \mid s \in S \wedge$ 
           $\text{RELEASE}([\text{node}_0(s', b) : s' \sim_{a \in A} s],$ 
           $\{ \}, \langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi)\}$ 
10:   return  $[\phi_0]_M$ 

```

Algorithm 5 Labelling $\langle\langle A^b \rangle\rangle \circ \phi$

```

1: function  $\text{NEXT}(B, C, \langle\langle A^b \rangle\rangle \circ \phi)$ 
2:   if  $B = [ ]$  then
3:     return true
4:    $n \leftarrow \text{hd}(B)$ 
5:    $\text{Act}_A \leftarrow \{\sigma \in D_A(s(n)) \mid \text{cons}(s(n), \sigma) \leq e(n) \wedge$ 
           $\text{out}(s(n), \sigma) \subseteq [\phi]_M \wedge \forall a \in A$ 
           $\text{if } \exists n' \in C : p(n) \cdot n \sim_a p(n')$ 
           $\text{then } \sigma_a = \text{act}_a(p(n')[1])\}$ 
6:   for  $\sigma \in \text{Act}_A$  do
7:     if  $\text{NEXT}(tl(B), C \cup \{\text{node}(n, \sigma, \text{hd}(\text{out}(s(n), \sigma)))\},$ 
           $\langle\langle A^b \rangle\rangle \circ \phi)$  then
8:       return true
9:   return false

```

LEMMA 3.1. *Algorithm 4 terminates in at most $O(|\phi| \times |M|^{k+1})$ steps, where k is the maximal value of the first resource bound in ϕ .*

PROOF. The algorithm for $\langle\langle A^b \rangle\rangle \circ \phi$ attempts to find an action which works (achieves ϕ) from all states indistinguishable from s by some agent in A . There are at most $|S|$ such states, and at most $|M|$ possible actions to try. In the worst case (when no action works in all states) we try every action in each state: $O(|M|^2)$ steps.

As before, the algorithms for $\langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$ and $\langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi$ are attempting to find a strategy of depth $\min_{a \in A}(b_{a1})$, but now from all

Algorithm 6 Labelling $\langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$

```

1: function  $\text{UNTIL}(B, C, \langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi)$ 
2:   if  $B = [ ]$  then
3:     return true
4:    $n \leftarrow \text{hd}(B)$ 
5:   if  $s(n) \in [\psi]_M$  then
6:     return  $\text{UNTIL}(tl(B), C \cup \{n\}, \langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi)$ 
7:   if  $s(n) \notin [\phi]_M$  then
8:     return false
9:    $\text{Act}_A \leftarrow \{\sigma \in D_A(s(n)) \mid \text{cons}(s(n), \sigma) \leq e(n) \wedge \forall a \in A$ 
           $\text{if } \exists n' \in C : p(n) \cdot n \sim_a p(n')[1, |p(n) \cdot n|]$ 
           $\text{then } \sigma_a = \text{act}_a(p(n')[|p(n) \cdot n| + 1])\}$ 
10:  for  $\sigma \in \text{Act}_A$  do
11:     $P \leftarrow \{\text{node}(n, \sigma, s') \mid s' \in \text{out}(s(n), \sigma)\}$ 
12:    if  $\text{UNTIL}(P \circ tl(B), C, \langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi)$  then
13:      return true
14:  return false

```

Algorithm 7 Labelling $\langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi$

```

1: function  $\text{RELEASE}(B, C, \langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi)$ 
2:   if  $B = [ ]$  then
3:     return true
4:    $n \leftarrow \text{hd}(B)$ 
5:   if  $s(n) \in [\psi]_M \cap [\phi]_M$  then
6:     return  $\text{RELEASE}(tl(B), C \cup \{n\}, \langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi)$ 
7:   if  $s(n) \notin [\psi]_M$  then
8:     return false
9:    $\text{Act}_A \leftarrow \{\sigma \in D_A(s(n)) \mid \forall a \in A$ 
           $\text{if } \exists n' \in C : p(n) \cdot n \sim_a p(n')[1, |p(n) \cdot n|]$ 
           $\text{then } \sigma_a = \text{act}_a(p(n')[|p(n) \cdot n| + 1])\}$ 
10:  for  $\sigma \in \text{Act}_A$  do
11:    if  $\text{cons}(s(n), \sigma) \not\leq e(n)$  then
12:       $n' \leftarrow \text{node}(n, \sigma, s' \in \text{out}(s(n), \sigma))$ 
13:      if  $\text{RELEASE}(tl(B), C \cup \{n'\}, \langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi)$  then
14:        return true
15:    else
16:       $P \leftarrow \{\text{node}(n, \sigma, s') \mid s' \in \text{out}(s(n), \sigma)\}$ 
17:      if  $\text{RELEASE}(P \circ tl(B), C, \langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi)$  then
18:        return true
19:  return false

```

indistinguishable states and satisfying additional constraints of uniformity. Considering all indistinguishable states adds an additional level (intuitively the root of the tree from which all indistinguishable initial states are reachable). Satisfying uniformity means having to backtrack to a successful subtree to try a different choice of actions even if the previous choice was successful (because the same choice does not work in an indistinguishable branch on another tree). In the worst case, we will consider all possible actions at each of $O(b)$ levels of the search tree. We repeat this for every subformula ($|\phi|$ many times). \square

LEMMA 3.2. *Algorithm 4 is correct.*

PROOF. We consider the cases of $\langle\langle A^b \rangle\rangle \circ \phi$, $\langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$ and $\langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi$.

The algorithm for $\langle\langle A^b \rangle\rangle \circ \phi$ places all states which are indistinguishable from the current state for one of the agents in A in the open list B . This ensures that a successful strategy (single action σ which is b -consistent and achieves ϕ) found in state s will be placed in the closed list C , and in states $s' \sim_a s$ (indistinguishable for the agent a) the same action σ_a will be attempted as part of the joint action σ' by A . If this does not result in a successful strategy in s' , the algorithm will backtrack and try another action for a in s . The algorithm returns true if and only if in all indistinguishable states, an action by A is found which always results in a state satisfying ϕ , is under the resource bound, and its a th component is the same in all \sim_a states. This guarantees that the algorithm found a one step strategy to satisfy the ϕ . In order to extend it to an arbitrary uniform strategy, we can simply select the first action in $D_a(s')$ for all sequences ending in s' and all $a \in A$. This will ensure that all a -indistinguishable sequences are assigned the same action.

$\langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$ implements the same idea as above, but with respect to multi-step strategies. Every time an action is selected on some path p , if $p' \sim_a p$ is in the closed list C , then a 's action after p is selected to be the same as that selected after p' . If this is not successful then eventually we will fail back to p' and try a different action there. If the algorithm returns true, then we are guaranteed that the strategy contained in C is uniform. We can easily extend the strategy contained in C to a uniform strategy, since we do not need to achieve any objectives after satisfying ψ .

$\langle\langle A^b \rangle\rangle \phi \mathcal{R} \psi$ is similar to $\langle\langle A^b \rangle\rangle \phi \mathcal{U} \psi$, but now we have an additional complication that actions selected to 'run out of resources' need to be in the closed list since they should also satisfy uniformity. This is ensured on lines 11-14 of the algorithm (we add a path ending with an 'expensive' action σ and an arbitrary successor n' to the closed list). \square

THEOREM 3.3. *The model-checking problem for $RB \pm ATL^\#$ with imperfect information and perfect recall is decidable in EXPSpace if the resource bounds are represented in unary.*

PROOF. In addition to the space required for the stack, we also need to store the closed list C . In the worst case, the closed list will contain all possible sequences of states of length at most $\max_{a \in A} (b_{a1})$, which is $O(|S|^k)$, where k is the maximal value of the first resource bound in ϕ . \square

4 RAL[#]

In this section we define a diminishing resource version of *resource agent logic* (RAL[#]) following [12], with modifications required for our setting (e.g., no infinite endowments).

The logic is defined over a set of agents Agt , a set of resource types Res , and a set of propositional symbols Π .

An *endowment (function)* $\eta : Agt \times Res \rightarrow \mathbb{N}$ assigns resources to agents; $\eta_a(r) = \eta(a, r)$ is the amount of resource agent a has of resource type r . En denotes the set of all possible endowments.

The formulas of RAL[#] are defined by:

$$\phi, \psi ::= p \mid \neg \phi \mid \phi \wedge \phi \mid \langle\langle A \rangle\rangle_B^\downarrow \circ \phi \mid \langle\langle A \rangle\rangle_B^\uparrow \circ \phi \mid \langle\langle A \rangle\rangle_B^\downarrow \phi \mathcal{U} \psi \mid \langle\langle A \rangle\rangle_B^\uparrow \phi \mathcal{U} \psi \mid \langle\langle A \rangle\rangle_B^\downarrow \phi \mathcal{R} \psi \mid \langle\langle A \rangle\rangle_B^\uparrow \phi \mathcal{R} \psi$$

where $p \in \Pi$ is a proposition, $A, B \subseteq Agt$ are sets of agents, and η is an endowment. A are called the proponents, and B the (resource-bounded) opponents.

Unlike in $RB \pm ATL^\#$, in RAL[#] there are two types of cooperation modalities, $\langle\langle A \rangle\rangle_B^\downarrow$ and $\langle\langle A \rangle\rangle_B^\uparrow$. In both types of cooperation modality, the actions performed by agents in $A \cup B$ consume and produce resources (actions by agents in $Agt \setminus (A \cup B)$ do not change their resource endowment). The meaning of $\langle\langle A \rangle\rangle_B^\uparrow \phi$ is otherwise the same as in $RB \pm ATL^\#$. The formula $\langle\langle A \rangle\rangle_B^\downarrow \phi$ on the other hand requires that the strategy uses the resources *currently* available to the agents.

The models of RAL[#] are resource-bounded concurrent game structures with diminishing resource (RB-CGS[#]). Strategies are also defined as for $RB \pm ATL^\#$. However, to evaluate formulas with a down arrow, such as $\langle\langle A \rangle\rangle_B^\downarrow \circ \phi$, we need the notion of *resource-extended computations*. A *resource-extended* computation $\lambda \in (S \times En)^+$ is a non-empty sequence over $S \times En$ such that the restriction to states (the first component), denoted by $\lambda|_S$, is a path in the underlying model. The projection of λ to the second component of each element in the sequence is denoted by $\lambda|_{En}$.

A (η, s_A, B) -computation is a resource-extended computation λ where for all $i = 1, \dots$ with $\lambda[i] := (s_i, \eta^i)$ there is an action profile $\sigma \in d(\lambda|_S[i])$ such that:

- (1) $\eta^0 = \eta$ (η describes the initial resource distribution);
- (2) $F_A(\lambda|_S[1, i]) = \sigma_A$ (A follow their strategy);
- (3) $\lambda|_S[i+1] = \delta(\lambda|_S[i], \sigma)$ (transition according to σ);
- (4) for all $a \in A \cup B$ and $r \in Res$: $\eta_a^i(r) \geq \text{cons}_r(\lambda|_S[i], \sigma_a)$ (each agent has enough resources to perform its action);
- (5) for all $a \in A \cup B$ and $r \in Res$: $\eta_a^{i+1}(r) = \eta_a^i(r) + \text{prod}_r(\lambda|_S[i], \sigma_a) - \text{cons}_r(\lambda|_S[i], \sigma_a)$ (resources are updated);
- (6) for all $a \in Agt \setminus (A \cup B)$ and $r \in Res$: $\eta_a^{i+1}(r) = \eta_a^i(r)$ (the resources of agents not in $A \cup B$ do not change).

The (η, B) -outcome of a strategy F_A in s , $out(s, \eta, F_A, B)$, is defined as the set of all (η, F_A, B) -computations starting in s . Truth is defined over a model M , a state $s \in S$, and an endowment η .

The semantics is given by the satisfaction relation \models where the cases for propositions, negation and conjunction are standard and omitted:

- $M, s, \eta \models \langle\langle A \rangle\rangle_B^\downarrow \circ \phi$ iff there is a strategy F_A for A such that for all $\lambda \in out(s, \eta, F_A, B)$, $|\lambda| > 1$ and $M, \lambda|_S[2], \lambda|_{En}[2] \models \phi$
- $M, s, \eta \models \langle\langle A \rangle\rangle_B^\uparrow \circ \phi$ iff there is a strategy F_A for A such that for all $\lambda \in out(s, \zeta, F_A, B)$, $|\lambda| > 1$ and $M, \lambda|_S[2], \lambda|_{En}[2] \models \phi$
- $M, s, \eta \models \langle\langle A \rangle\rangle_B^\downarrow \phi \mathcal{U} \psi$ iff there is a strategy F_A for A such that for all $\lambda \in out(s, \eta, F_A, B)$, there exists i with $1 \leq i \leq |\lambda|$ and $M, \lambda|_S[i], \lambda|_{En}[i] \models \psi$ and for all j with $1 \leq j < i$, $M, \lambda|_S[j], \lambda|_{En}[j] \models \phi$
- $M, s, \eta \models \langle\langle A \rangle\rangle_B^\uparrow \phi \mathcal{U} \psi$ iff there is a strategy F_A for A such that for all $\lambda \in out(s, \zeta, F_A, B)$, there exists i with $1 \leq i \leq |\lambda|$ and $M, \lambda|_S[i], \lambda|_{En}[i] \models \psi$ and for all j with $1 \leq j < i$, $M, \lambda|_S[j], \lambda|_{En}[j] \models \phi$
- $M, s, \eta \models \langle\langle A \rangle\rangle_B^\downarrow \phi \mathcal{R} \psi$ iff there is a strategy F_A for A such that for all $\lambda \in out(s, \eta, F_A, B)$, either there exists i with $1 \leq i \leq |\lambda|$ and $M, \lambda|_S[i], \lambda|_{En}[i] \models \psi \wedge \phi$ and for all j with $1 \leq j < i$, $M, \lambda|_S[j], \lambda|_{En}[j] \models \psi$; or, for all j with $1 \leq j \leq |\lambda|$, $M, \lambda|_S[j], \lambda|_{En}[j] \models \psi$

$M, s, \eta \models \langle\langle A \rangle\rangle_B^{\zeta} \phi \mathcal{R} \psi$ iff there is a strategy F_A for A such that for all $\lambda \in \text{out}(s, \zeta, F_A, B)$, either there exists i with $1 \leq i \leq |\lambda|$ and $M, \lambda|_S[i], \lambda|_{En}[i] \models \psi \wedge \phi$ and for all j with $1 \leq j < i$, $M, \lambda|_S[j], \lambda|_{En}[j] \models \psi$; or, for all j with $1 \leq j \leq |\lambda|$, $M, \lambda|_S[j], \lambda|_{En}[j] \models \psi$

The model checking algorithms for $\text{RAL}^\#$ are similar to those given for $\text{RB} \pm \text{ATL}^\#$ in Section 2 in that they proceed by and-or depth first search. However, in this case, the nodes in the search tree also include information about the current proponent and (resource-bounded) opponent coalitions, and the functions that construct nodes are redefined as $\text{node}_0(s, b, A, B)$ and $\text{node}(n, \sigma, s', A, B)$ where A are the proponents and B are the resource-bounded opponents.

The model checking algorithm for $\text{RAL}^\#$ is shown in Algorithm 8, and takes as input a model M , a formula ϕ , and an initial endowment η , and labels the set of states $[\phi]_M^\eta$, where $[\phi]_M^\eta = \{s \mid M, s, \eta \models \phi\}$ is the set of states satisfying ϕ . $\text{RAL}^\#$ -LABEL simply calls the function STRATEGY to label states with ϕ . pr and op are functions that return the proponents $A \subseteq \text{Agt}$ and the resource-bounded opponents $B \subseteq \text{Agt}$ respectively if ϕ is of the form $\langle\langle A \rangle\rangle_B^* \phi$, $\langle\langle A \rangle\rangle_B^* \psi_1 \mathcal{U} \psi_2$, $\langle\langle A \rangle\rangle_B^* \psi_1 \mathcal{R} \psi_2$ where $*$ is either \downarrow or an endowment, or \emptyset otherwise.

Algorithm 8 Labelling ϕ

```

1: procedure  $\text{RAL}^\#$ -LABEL( $M, \phi, \eta$ )
2:    $[\phi]_M^\eta \leftarrow \{q \mid q \in S \wedge$ 
      STRATEGY( $\text{node}_0(q, \eta, pr(\phi), op(\phi)), \phi)\}$ 

```

The function STRATEGY is shown in Algorithm 9 and proceeds by depth-first and-or search. We process each coalition modality in turn, starting from the outermost modality. The logical connectives are standard, and simply call STRATEGY on the subformulas. Each temporal operator is handled by a separate function: NEXT for $\bigcirc \psi$, UNTIL for $\phi \mathcal{U} \psi$, and RELEASE for $\phi \mathcal{R} \psi$.

Algorithm 9 Strategy

```

1: function STRATEGY( $n, \phi$ )
2:   case  $\phi = p \in \Pi$ 
3:     return  $s(n) \in \pi(p)$ 
4:   case  $\phi = \neg \psi$ 
5:     return  $\neg$ STRATEGY( $\text{node}_0(s(n), e(n), pr(n), op(n)), \psi$ )
6:   case  $\phi = \psi_1 \vee \psi_2$ 
7:     return STRATEGY( $\text{node}_0(s(n), e(n), pr(n), op(n)), \psi_1$ )  $\vee$ 
      STRATEGY( $\text{node}_0(s(n), e(n), pr(n), op(n)), \psi_2$ )
8:   case  $\phi = \langle\langle A \rangle\rangle_B^{\downarrow} \phi$ 
9:     return NEXT( $\text{node}_0(s(n), e(n), A, B), \phi$ )
10:  case  $\phi = \langle\langle A \rangle\rangle_B^{\zeta} \phi$ 
11:    return NEXT( $\text{node}_0(s(n), \zeta, A, B), \phi$ )
12:  case  $\phi = \langle\langle A \rangle\rangle_B^{\downarrow} \psi_1 \mathcal{U} \psi_2$ 
13:    return UNTIL( $\text{node}_0(s(n), e(n), A, B), \phi$ )
14:  case  $\phi = \langle\langle A \rangle\rangle_B^{\zeta} \psi_1 \mathcal{U} \psi_2$ 
15:    return UNTIL( $\text{node}_0(s(n), \zeta, A, B), \phi$ )
16:  case  $\phi = \langle\langle A \rangle\rangle_B^{\downarrow} \psi_1 \mathcal{R} \psi_2$ 
17:    return RELEASE( $\text{node}_0(s(n), e(n), A, B), \phi$ )
18:  case  $\phi = \langle\langle A \rangle\rangle_B^{\zeta} \psi_1 \mathcal{R} \psi_2$ 
19:    return RELEASE( $\text{node}_0(s(n), \zeta, A, B), \phi$ )

```

The function NEXT for formulas of types $\langle\langle A \rangle\rangle_B^{\downarrow} \phi$ and $\langle\langle A \rangle\rangle_B^{\zeta} \phi$ is shown in Algorithm 10 and is straightforward. We simply check if there is an action of A that is possible given the current endowment (lines 2–4), and where in all outcome states A has a strategy to enforce ϕ (lines 6–10). Note that the recursive call (line 8) is to STRATEGY, to correctly determine the endowments for the new search in both the case where ϕ specifies a fresh endowment or the resources currently available to the agents (i.e., down arrow).

Algorithm 10 Next (both types of modalities)

```

1: function NEXT( $n, \langle\langle A \rangle\rangle_B^* \phi$ )
2:    $Act_A \leftarrow \{\sigma' \in D_A(s(n)) \mid \text{cons}(\sigma') \leq e_A(n)\}$ 
3:   for  $\sigma' \in Act_A$  do
4:      $Act_{Agt} \leftarrow \{\sigma \in D(s(n)) \mid \sigma_A = \sigma' \wedge$ 
      cons( $\sigma_B$ )  $\leq e_B(n)\}$ 
5:      $strat \leftarrow true$ 
6:     for  $\sigma \in Act_{Agt}$  do
7:        $s' \leftarrow \delta(s(n), \sigma)$ 
8:        $strat \leftarrow strat \wedge$  STRATEGY( $\text{node}(n, \sigma, s', A, B), \phi$ )
9:     if  $strat$  then
10:      return true
11:   return false

```

The function UNTIL for formulas of types $\langle\langle A \rangle\rangle_B^{\downarrow} \phi \mathcal{U} \psi$ and $\langle\langle A \rangle\rangle_B^{\zeta} \phi \mathcal{U} \psi$ is shown in Algorithm 11. If A have a strategy to enforce ψ , we return true (lines 2–3). We then check if it is possible to enforce ϕ in n , and terminate the search with false if it is not (lines 4–5). Otherwise the search continues. Each action available at $s(n)$ is considered in turn (lines 6–14). For each action $\sigma' \in Act_A$, we check whether a recursive call of the algorithm returns true in all outcome states s' of σ' (i.e., σ' is part of a successful strategy). If such a σ' is found, the algorithm returns true. Otherwise the algorithm returns false.

The function RELEASE for formulas of types $\langle\langle A \rangle\rangle_B^{\downarrow} \phi \mathcal{R} \psi$ and $\langle\langle A \rangle\rangle_B^{\zeta} \phi \mathcal{R} \psi$ is similar (see Algorithm 12).

Algorithm 11 Until (both types of modalities)

```

1: function UNTIL( $n, \langle\langle A \rangle\rangle_B^* \phi \mathcal{U} \psi$ )
2:   if STRATEGY( $n, \psi$ ) then
3:     return true
4:   if  $\neg$ STRATEGY( $n, \phi$ ) then
5:     return false
6:    $Act_A \leftarrow \{\sigma' \in D_A(s(n)) \mid \text{cons}(\sigma') \leq e_A(n)\}$ 
7:   for  $\sigma' \in Act_A$  do
8:      $Act_{Agt} \leftarrow \{\sigma \in D(s(n)) \mid \sigma_A = \sigma' \wedge$ 
      cons( $\sigma_B$ )  $\leq e_B(n)\}$ 
9:      $strat \leftarrow true$ 
10:    for  $\sigma \in Act_{Agt}$  do
11:       $s' \leftarrow \delta(s(n), \sigma)$ 
12:       $strat \leftarrow strat \wedge$ 
      UNTIL( $\text{node}(n, \sigma, s', A, B), \langle\langle A \rangle\rangle_B^* \phi \mathcal{U} \psi$ )
13:    if  $strat$  then
14:      return true
15:   return false

```

LEMMA 4.1. Algorithm 9 terminates in $O(|M|^{|\phi|})$ steps, where the bounds in ϕ are written in unary.

Algorithm 12 Release (both types of modalities)

```

1: function RELEASE( $n, \langle\langle A \rangle\rangle_B^* \phi \mathcal{R} \psi$ )
2:   if  $\neg$ STRATEGY( $n, \psi$ ) then
3:     return false
4:   if STRATEGY( $n, \phi$ ) then
5:     return true
6:   if  $\exists \sigma \in D_A$  s.t.  $\text{cons}(s(n), \sigma) \not\leq e_A(n)$  then
7:     return true
8:    $Act_A \leftarrow \{\sigma' \in D_A(s(n)) \mid \text{cons}(\sigma') \leq e_A(n)\}$ 
9:   for  $\sigma' \in Act_A$  do
10:     $Act_{Agt} \leftarrow \{\sigma \in D(s(n)) \mid \sigma_A = \sigma' \wedge$ 
11:       $\text{cons}(\sigma_B) \leq e_B(n)\}$ 
12:     $strat \leftarrow true$ 
13:    for  $\sigma \in Act_{Agt}$  do
14:       $s' \leftarrow \delta(s(n), \sigma)$ 
15:       $strat \leftarrow strat \wedge$ 
16:        RELEASE( $node(n, \sigma, s', A, B), \langle\langle A \rangle\rangle_B^* \phi \mathcal{R} \psi$ )
17:    if  $strat$  then
18:      return true
19:  return false
    
```

PROOF. The only difference between the $RAL^\#$ algorithms and the algorithms in section 2 is the fact that in the case of $RAL^\#$ we cannot label states with subformulas. For example, we cannot find states satisfying $\langle\langle A \rangle\rangle_B^\downarrow \phi \mathcal{U} \psi$ because we do not know which endowment the \downarrow refers to. When verifying a formula with non-propositional subformulas, for example $\langle\langle A \rangle\rangle_B^\downarrow \phi \mathcal{U} \psi$ again, where ϕ and ψ are not propositional, we have to make recursive calls to check whether the current state satisfies ϕ or ψ with the current endowment. Hence the checks for STRATEGY(n, ϕ) instead of checking whether $s(n) \in [\phi]_M$. However the recursive calls are always to formulas of lower complexity, and it is easy to show that in the propositional case they do terminate, and that under the inductive assumption if lower complexity calls terminate, then the calls to $\langle\langle A \rangle\rangle_B^* \phi$, $\langle\langle A \rangle\rangle_B^* \phi \mathcal{U} \psi$ and $\langle\langle A \rangle\rangle_B^* \phi \mathcal{R} \psi$ terminate.

The algorithm again performs depth first and-or search, but now up to the depth determined by the nestings of modalities in ϕ : we need to take the sum of the minimal bounds for the first resource occurring in the endowment of some resource bounded agent in nested formulas to find the maximal depth of the tree. We can ignore \downarrow endowments because they will use the amount of the first resource remaining from the outer modalities. \square

LEMMA 4.2. *Algorithm 9 is correct.*

PROOF. Assuming that calls to STRATEGY(n, ϕ) terminate and have the same effect as checking whether $s(n) \in [\phi]_M$, the algorithms are the same as for $RB \pm ATL^\#$. The only small difference is that we remember the current endowment and pass it to the \downarrow modalities as if it was an explicit bound b in $RB \pm ATL^\#$. \square

THEOREM 4.3. *The model-checking problem for $RAL^\#$ is decidable in PSPACE (if resource bounds are written in unary).*

PROOF. From the two lemmas above it follows that Algorithm 9 is a terminating and correct model-checking algorithm for $RAL^\#$. The space it is using on the stack is polynomial in the size of the formula (it is the sum of nested resource bounds on the first resource

for the minimally endowed agents). After at most $O(k)$ steps, where k is the maximal value of the first resource bound in ϕ , the endowment becomes negative for one of the agents, and the algorithm terminates. \square

5 CONCLUSION

In this paper we studied resource logics over models with a diminishing resource. We gave new and simple model-checking algorithms for the versions of $RB \pm ATL$, $RB \pm ATL_{iR}$ and RAL with a diminishing resource. We believe that settings where one of the resources is always consumed are quite common, and our results may therefore be of practical interest. It was known that the model checking problem for $RB \pm ATL$ is decidable, but our complexity result for $RB \pm ATL^\#$ is new. Decidability of the model checking problem for RAL follows from a more general result on bounded models from [12], but no model checking algorithm was given there. The model checking algorithm for $RAL^\#$ is different from the algorithm for the decidable fragment of RAL presented in [3] because it works for the full RAL rather than just for the positive fragment of proponent-restricted RAL in [3].

REFERENCES

- [1] Parosh Aziz Abdulla, Richard Mayr, Arnaud Sangnier, and Jeremy Sproston. 2013. Solving Parity Games on Integer Vectors. In *CONCUR'13*, Vol. 8052. Springer, 106–120.
- [2] Natasha Alechina, Nils Bulling, Stéphane Demri, and Brian Logan. 2016. On the Complexity of Resource-Bounded Logics. In *Proceedings of the 10th International Workshop on Reachability Problems (RP 2016) (Lecture Notes in Computer Science)*, Kim Guldstrand Larsen, Igor Potapov, and Jiri Srba (Eds.), Vol. 9899. Springer, 36–50.
- [3] Natasha Alechina, Nils Bulling, Brian Logan, and Hoang Nga Nguyen. 2017. The virtues of idleness: A decidable fragment of resource agent logic. *Artificial Intelligence* 245 (2017), 56–85. <https://doi.org/10.1016/j.artint.2016.12.005>
- [4] Natasha Alechina, Brian Logan, Hoang Nga Nguyen, and Franco Raimondi. 2014. Decidable Model-Checking for a Resource Logic with Production of Resources. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI-2014)*, Torsten Schaub, Gerhard Friedrich, and Barry O’Sullivan (Eds.). IOS Press, 9–14.
- [5] Natasha Alechina, Brian Logan, Hoang Nga Nguyen, and Franco Raimondi. 2017. Model-checking for Resource-Bounded ATL with production and consumption of resources. *J. Comput. System Sci.* 88 (September 2017), 126–144. <https://doi.org/doi.org/10.1016/j.jcss.2017.03.008>
- [6] Natasha Alechina, Brian Logan, Hoang Nga Nguyen, and Abdur Rakib. 2009. A Logic for Coalitions with Bounded Resources. In *Proceedings of the Twenty First International Joint Conference on Artificial Intelligence (IJCAI 2009)*, Craig Boutilier (Ed.), Vol. 2. AAAI Press, 659–664.
- [7] Natasha Alechina, Brian Logan, Hoang Nga Nguyen, and Abdur Rakib. 2010. Resource-bounded alternating-time temporal logic. In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, Wiebe van der Hoek, Gal Kaminka, Yves Lespérance, Michael Luck, and Sandip Sen (Eds.). IFAAMAS, 481–488.
- [8] Natasha Alechina, Brian Logan, Hoang Nga Nguyen, and Abdur Rakib. 2011. Logic for coalitions with bounded resources. *Journal of Logic and Computation* 21, 6 (2011), 907–937.
- [9] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. 2002. Alternating-Time Temporal Logic. *Journal of the ACM* 49, 5 (2002), 672–713.
- [10] Tomas Brázdil, Petr Jancar, and Antonín Kucera. 2010. Reachability Games on Extended Vector Addition Systems with States. In *ICALP'10*, Vol. 6199. Springer, 478–489.
- [11] Nils Bulling and Berndt Farwer. 2010. Expressing Properties of Resource-Bounded Systems: The Logics RBTL and RBTL*. In *Computational Logic in Multi-Agent Systems - 10th International Workshop, CLIMA X, Revised Selected and Invited Papers (Lecture Notes in Computer Science)*, Vol. 6214. 22–45.
- [12] Nils Bulling and Berndt Farwer. 2010. On the (Un-)Decidability of Model Checking Resource-Bounded Agents. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI 2010)*, Helder Coelho, Rudi Studer, and Michael Wooldridge (Eds.). IOS Press, 567–572.
- [13] Nils Bulling and Valentin Goranko. 2013. How to Be Both Rich and Happy: Combining Quantitative and Qualitative Strategic Reasoning about Multi-Player Games (Extended Abstract). In *Proceedings of the 1st International Workshop*

- on *Strategic Reasoning (SR 2013) (Electronic Proceedings in Theoretical Computer Science)*, Fabio Mogavero, Aniello Murano, and Moshe Y. Vardi (Eds.), Vol. 112. 33–41.
- [14] Jean-Baptiste Courtois and Sylvain Schmitz. 2014. Alternating Vector Addition Systems with States. In *MFCS'14*, Vol. 8634. Springer, 220–231.
- [15] Aldric Degorre, Laurent Doyen, Raffaella Gentilini, Jean-François Raskin, and Szymon Toruńczyk. 2010. Energy and Mean-Payoff Games with Imperfect Information. In *Computer Science Logic, 24th International Workshop, CSL 2010, 19th Annual Conference of the EACSL, Brno, Czech Republic, August 23-27, 2010. Proceedings (Lecture Notes in Computer Science)*, Anuj Dawar and Helmut Veith (Eds.), Vol. 6247. Springer, 260–274. <https://doi.org/10.1007/978-3-642-15205-4>
- [16] Dario Della Monica and Giacomo Lenzi. 2012. On a Priced Resource-bounded Alternating μ -Calculus. In *Proceedings of the 4th International Conference on Agents and Artificial Intelligence (ICAART 2012)*, Joaquim Filipe and Ana L. N. Fred (Eds.). SciTePress, 222–227.
- [17] Dario Della Monica, Margherita Napoli, and Mimmo Parente. 2011. On a Logic for Coalitional Games with Priced-Resource Agents. *Electronic Notes in Theoretical Computer Science* 278 (2011), 215–228.
- [18] Dario Della Monica, Margherita Napoli, and Mimmo Parente. 2013. Model checking coalitional games in shortage resource scenarios. In *Proceedings of the Fourth International Symposium on Games, Automata, Logics and Formal Verification (GandALF 2013) (Electronic Proceedings in Theoretical Computer Science)*, Gabriele Puppis and Tiziano Villa (Eds.), Vol. 119. 240–255.
- [19] Catalin Dima and Ferucio Laurentiu Tiplea. 2011. Model-checking ATL under Imperfect Information and Perfect Recall Semantics is Undecidable. *CoRR* abs/1102.4225 (2011). <http://arxiv.org/abs/1102.4225>
- [20] Marcin Jurdzinski, Ranko Lazic, and Sylvain Schmitz. 2015. Fixed-Dimensional Energy Games are in Pseudo-Polynomial Time. In *ICALP'15*, Vol. 9135. Springer, 260–272.
- [21] François Laroussinie, Nicolas Markey, and Ghassan Oreiby. 2008. On the Expressiveness and Complexity of ATL. *Logical Methods in Computer Science* 4, 2 (2008). [https://doi.org/10.2168/LMCS-4\(2:7\)2008](https://doi.org/10.2168/LMCS-4(2:7)2008)
- [22] Bastien Maubert and Sophie Pinchinat. 2014. A General Notion of Uniform Strategies. *IGTR* 16, 1 (2014). <https://doi.org/10.1142/S0219198914400040>
- [23] Steen Vester. 2013. Alternating-time temporal logic with finite-memory strategies. In *Proceedings Fourth International Symposium on Games, Automata, Logics and Formal Verification, GandALF 2013, Borca di Cadore, Dolomites, Italy, 29-31th August 2013 (EPTCS)*, Gabriele Puppis and Tiziano Villa (Eds.), Vol. 119. 194–207. <https://doi.org/10.4204/EPTCS.119>